
TCP Performance Implications of Network Asymmetry

draft-ietf-pilc-asym-03.txt

Authors:

Hari Balakrishnan
Venkata N. Padmanabhan
Gorry Fairhurst
Mahesh Sooriyabandara

Revisions to the draft

Draft -00 (Sept 99)

Issued in response to WG charter

Draft -01

Added issues and techniques

Draft -02

Minor corrections

Draft -03

Major Revision

- Added examples of asymmetric networks

- Added techniques

- Techniques organised by type

- Identify techniques in use


- Relationship to PEP clarified

- Recommendations added



Bandwidth Asymmetry

Downstream/Forward: Constrained throughput



Upstream/Reverse: Limited ACK rate

Most applications send more than they receive

High asymmetry causes return pipe to fill with ACKs

Asymmetry

Examples of Asymmetry

Asked WG for examples
- Please do send more!

Links which benefit from a lower return rate


Shared medium access
High per packet “cost” for many “radio” links

Asymmetry benefits such links benefit *by design*

Important note:

These slides use a satellite example
- The same applies for the other subnetworks !!!


Implications



- (i) **ACK rate controls TCP send rate (self-clocking)**
 - cwnd* opens slowly (low ACK rate)
 - Cumulative ACKs generate TCP DATA bursts
- (ii) **ACK Queue builds**
 - May drop ACKs
 - Increasing RTT, TCP RTO may expire *before* loss
 - Slowed reaction time of protocol, (FR etc)

Types of Mitigations

(a) End-to-End (various)



(b) Transparent (4 types)

End-to-End Mitigations

Modified Delayed ACKs

REC: Don't use (difficult to select d)

Large MSS

REC: Don't use IP fragmentation

Dynamically vary d

REC: Don't use - remain a research area

Other TCP Sender Modifications

REC: Don't use

Transparent Mitigations (type 0)

Header Compression

Reduces size of ACK

RFC1144 (V-J HC)

REC: Widely implemented and used

May use if low error rate, ordered delivery

Benefit with low-to-moderate asymmetry

Robust Header Compression (See ROHC WG)


REC: Benefit with low-to-moderate asymmetry

Benefit with low-to-moderate asymmetry

Does not reduce ACK rate

Does not mitigate with upstream DATA

Transparent Mitigations (type 1)



Techniques applied before the upstream bottleneck

ACK Filtering / Suppression

REC: Major benefit, has been deployed
May lead to TCP bursts

ACK Decimation

REC: Major benefit, has been deployed
May lead to TCP bursts
Some inelegant recovery

Transparent Mitigations (type 2)

Techniques applied after the upstream bottleneck
Mitigates the effect of *stretch ACKs* (TCP DATA bursts)

ACK Reconstruction (implicit)

REC: Desirable

Appropriate algorithms remain a research issue


ACK Compacting / Comanding (explicit)

REC: Desirable

Appropriate algorithms remain a research issue

Are security recommendations (packet amplifier) enough?

Shared Reverse (uplink)



Sharing reduces capacity per flow for uplink ACKs


- (i) ACKs from multiple flows
- (ii) DATA sharing with ACKs

Prone to ACK Compression

Often a KEY FACTOR

Transparent Mitigations (type 3)

Reverse link scheduling
Mitigate effect of sharing



Per-Flow queues

REC: Widely implemented
Desirable for all low speed links

ACKs First Scheduling

Separate queues for DATA and ACKs (hi priority)
Used with a scheme to reduce ACK rate

REC: Promising
Appropriate algorithms remain a research issue

Conclusions



Major revision (-03) completed

Thanks to ALL who provided new input

Intention to correct known mistakes (-04) April

More Comments VERY Welcome:

Taxonomy correct?

More example networks?

More mitigations?

RECOMMENDATIONS CORRECT?